

SYSTEM AND METHOD FOR CONNECTIONLESS/CONNECTION
ORIENTED SIGNAL TRANSPORT

RELATED APPLICATION

This application claims the benefit of U.S. Provisional Application Serial No. 60/202,190, entitled Internaet Protocol Transport, filed May 5, 2000, which is hereby incorporated by reference.

TECHNICAL FIELD OF THE INVENTION

The present invention relates to the field of communication systems, and more particularly to a system and method for communicating connectionless and/or connection oriented signals using one or more common network elements.

BACKGROUND OF THE INVENTION

Existing networks typically operate in one of two modes, connection oriented mode or connectionless mode. In a connection oriented mode of operation, signals are communicated over specified paths from a source network element to a destination network element. Multi-Protocol Label Switching (MPLS) signals with, for example, Asynchronous Transport Mode (ATM), Frame Relay, or packet-over-SONET encoding are examples of connection oriented signals. In a connectionless mode of operation, signals are communicated with less regard for the particular path traversed between source and destination network elements. Connectionless signaling typically focuses on the destination address, or other identification, rather than any particular path between source and destination network elements. Internet Protocol (IP), IPX, and SNA packet switching are examples of connectionless signal transport.

Connection oriented and connectionless communication each have their virtues. Due to the tracking of per-connection information in a connection oriented network, the connection oriented network provides efficient bandwidth management, which enables quality of service (Qos) aware services and traffic engineering ability. Connectionless networks, while generally less capable of providing quality of service aware services and traffic engineering capabilities, provide advantages over connection oriented networks in terms of their simplicity, reliability, and scalability.

For example, connectionless networks do not generally require per-connection maintenance constructions to ensure complete state restoration upon recovery from a node failure. Rather, recovery from a nodal failure in a connectionless network typically only involves

resynchronizing the routing database. In terms of scaleability, the connectionless network elements need only track the other network elements in the network, while the connection oriented network elements generally track each 5 flow between various network elements. Because more than one flow can exist between network elements, the scaleability of connection oriented networks is generally limited if the network is to properly track various flow information. The lack of scaleability of conventional 10 connection oriented networks has generally limited their applicability to geographically limited networks.

These and other network differences have generally resulted in connection oriented signals being transported over one set of network elements, while 15 connectionless signals are communicated using another set of network elements. As a result, conventional network architectures, particularly wide area architectures have not supported both connectionless and connection oriented signaling.

20

SUMMARY OF THE INVENTION

In accordance with the present invention, a method of communicating connectionless and connection oriented signals using at least one common network element comprises 5 receiving connectionless and connection oriented signals from a plurality of source peripheral network elements and determining a signaling type associated with each signal, the signaling type comprising connectionless signaling or connection oriented signaling. The method further includes 10 appending a transport label to each signal, each transport label comprising an indication of the signal's signaling type, and communicating the signals and appended transport labels toward destination peripheral network elements according to signaling procedures associated with each 15 signal's signaling type.

Various embodiments of the present invention may exhibit some, none, or all of the following technical advantages. For example, the invention provides an 20 architecture that captures the advantages of connectionless and connection oriented signaling in a single network. The invention provides the ability to perform traffic engineering and bandwidth management functions associated with connection oriented service, while facilitating a scalable, geographically diverse network topology, which is 25 simple and reliable.

Other technical advantages are readily apparent to one of skill in the art from the attached figures, description, and claims.

BRIEF DESCRIPTION OF THE DRAWINGS

For a more complete understanding of the present invention, and for further features and advantages thereof, reference is now made to the following description taken in conjunction with the accompanying drawings, in which:

FIGURE 1 is a block diagram of an exemplary system operable to facilitate communication of connection oriented signals and connectionless signals using common network elements;

FIGURE 2 is a block diagram of an exemplary network element constructed according to the teachings of the present invention;

FIGURE 3A is a block diagram of at least a portion of one exemplary embodiment of a signal format constructed according to the teachings of the present invention;

FIGURE 3B is a block diagram of one exemplary embodiment of a transport label constructed according to the teachings of the present invention;

FIGURE 4 is a block diagram of an exemplary format look-up table constructed according to the teachings of the present invention;

FIGURE 5 is a block diagram of an exemplary signal format constructed according to the teachings of the present invention;

FIGURE 6A-6E are block diagrams of exemplary portions of signals communicated using the system shown in FIGURE 1; and

FIGURE 7 is a flow chart showing an exemplary method of communicating connectionless and connection oriented signals over common network elements.

DETAILED DESCRIPTION OF THE INVENTION

FIGURE 1 is a block diagram of an exemplary system 10 operable to facilitate communication of connection oriented signals and connectionless signals using common network elements. Throughout this description, the term "connection oriented signal" refers to a signal associated with a particular path, or portion of a path, from a source network element to a destination network element. In connection oriented signaling, the signal's path through the network elements is generally determined before the signal's payload traverses any of the network elements. Multi-Protocol Label Switching (MPLS) signals with ATM, Frame Relay, or Packet-over SONET encoding are examples of connection oriented signals.

Throughout this description, the term "connectionless signal" refers to a signal that is not necessarily associated with any particular path from a source network element to a destination network element. In connectionless signaling, routing determinations, such as determining the next network element in the path, are generally made at each node in the path, where each node independently identifies the next node in the forwarding chain. IP and IPx signals are examples of connectionless signals.

System 10 includes a core cloud 12 comprising one or more core network elements 14a-14n (referred to generally as "core network elements 14"). Core network elements 14 can communicate with one another via communication links 16a-16n, and with one or more peripheral network elements 18-24 via communication links 30a-30n. Communication links 16 and 30 may comprise any wireless, wireline, fiber, or other communication medium or combinations of media. Signals communicated via communication links 16 and/or 30 may be electrical, optical, or any combination thereof.

Peripheral network elements 18-24 facilitates communication between core cloud 12 and other network elements, and/or other networks, such as networks 32, 34, 36, and 38. In the illustrated embodiment, peripheral network elements 18-24 comprise routers. Peripheral network elements could also comprise other network elements, such as next-generation access devices. In the illustrated embodiment, peripheral network element 18 comprises a label switching router (LSR). Throughout this document, the term "label switching router" refers to hardware operable to route MPLS or other label-switched signals. Label switching router 18 couples to network 32 via a communication link 50. Similarly, label switching router 22 couples to network 36 via a communication link 54. Throughout this document, the term "couple" refers to a direct or indirect communication between two or more elements. Elements said to be "coupled" to one another may, but need not, be physically connected.

Label switching routers 18 and 22 facilitate routing functions for signals originated or forwarded by interface equipment 40a-40n and 44a-44n / 46a-46n, respectively, and communicated over networks 32 and 36, respectively. Interface units 40, 44, and 46 may be, for example, personal computers, servers, switches, routers, or any other network equipment operable to originate and/or forward communication signals. In the illustrated embodiment, interface units 40, 44, and 46 operate to communicate MPLS or other label switched signals with ATM, Frame Relay, and/or packet-over-SONET encoding. Other interface units communicating other types of signals could also couple to networks 32 and 36 without departing from the scope of the invention.

In the illustrated embodiment, peripheral network elements 20 and 24 comprise routers (R) operable to route

connectionless signals, such as Internet Protocol (IP) signals. Routers 20 and 24 facilitate routing functions for signals originated or forwarded by interface equipment 42a-42n and 48a-48n, respectively, and communicated over networks 34 and 38, respectively. Interface units 42 and 48 may be, for example, personal computers, servers, switches, routers, or any other network equipment operable to originate and/or forward communication signals to routers 20 and 24 through networks 34 and 38, respectively.

Networks 32, 34, 36, and 38 may comprise any suitable wireline or wireless system that supports communication between network elements using ground-based and/or space-based components. For example, networks 32-38 may be public switched telephone networks (PSTN), integrated services digital networks (ISDN), local area networks (LAN), metropolitan area networks (MAN), wide area networks (WAN), or other communication systems or combination of communication systems at one or more locations. Each of networks 32-38 may comprise a single network, or multiple separate networks. In addition, although for simplicity each network 32, 34, 36, and 38 has been shown as coupled to only one router, it should be understood that networks 32-38 could couple to or include any number of routers, switches, personal computers, servers, or other communication and/or computing devices.

In operation, core cloud 12 receives both connectionless and connection oriented signals from peripheral network elements 18-24, and routes those signals through core network 12 to another appropriate peripheral network element 18-24 according to routing procedures associated with the signal type received. In a particular embodiment, an ingress core network element 14 receives an incoming signal from a peripheral network element 18-24, which is bound for another peripheral network element

18-24. The ingress core network element 14 determines whether the signal is to be transported using connectionless or connection oriented signaling - or a combination of connectionless and connection oriented 5 signaling - and appends a transport label to the incoming signal. The transport label identifies the type of signaling the incoming signal will experience as it traverses core cloud 12.

The ingress core network element 14 identifies an 10 egress core network element 14 associated with a destination peripheral network element 18-24, and communicates the signal toward egress core network element 14. Depending on the type of signal, ingress core network element 14 may route the signal using connectionless 15 routing, connection oriented routing, or a combination of the two. Core network elements 14 residing between the ingress core network element 14 and the egress core network element 14 receive the signal with an appended transport label, and process the signal in accordance with the 20 transport label. The egress core network element 14, receives the signal and returns the signal to its original format by, for example, removing the appended transport label, and communicates the signal to the destination peripheral network element 18-24.

25 System 10 provides an architecture that captures the advantages of connectionless and connection oriented signaling in a single network. For example, system 10 provides the ability to perform traffic engineering and bandwidth management functions associated with connection 30 oriented service. At the same time, using the same core network elements, system 10 facilitates a scalable, geographically diverse network topology, which is simple and reliable.

FIGURE 2 is a block diagram of an exemplary core network element 14 constructed according to the teachings of the present invention. Each core network element 14 includes a core interface 110 operable to facilitate communication between various core network elements 14. Core network elements 14 which communicate with one or more peripheral network elements 18-24 may also include a peripheral interface (periph. I/F) 112, operable to facilitate communication between a core network element 14 and one or more associated peripheral network elements 18-24. Core interface 110 and peripheral interface 112 include the appropriate hardware, software, and/or firmware to communicate using a direct connection to a PSTN or ISDN, a connection through a LAN, WAN, or global computer network such as the Internet, or any other suitable communication link that allows core network elements 14 to communicate with other core network elements 14 and/or peripheral network elements 18-24.

A processor 114 communicates with core interface 110 and peripheral interface 112. Throughout this document, the term "processor" refers to hardware, software, firmware, or any combination thereof in one or more locations operable to provide signal and/or data processing functionality. In the illustrated embodiment, processor 114 comprises a portion of an application-specific integrated circuit (ASIC). Other types of processors could be used.

Although the illustrated embodiment shows processor 114 and interfaces 110, 112 as separate entities, some or all of the functionality of either or both of interfaces 110 and 112 could be implemented in processor 114 without departing from the scope of the present invention. Processor 114 also communicates with a memory 116. Memory 116 may comprise any storage device or combination of

storage devices accessible to processor 114. Memory 116 may comprise an integral part of core network element 14, or may reside remotely from and be accessible to core network element 14. In the illustrated embodiment, 5 memory 116 includes a plurality of look-up tables for use in processing and/or routing signals received at core network element 14. Throughout this document, the term "table" refers to any data structure, compilation, or arrangement of information operable to facilitate storage, 10 retrieval, and/or manipulation of that information. In the illustrated embodiment, memory 116 includes a node identification (Node ID) look-up table 118. Node ID look-up table 118 includes information, such as, network addresses for some or all of core network elements 14 in 15 core cloud 12.

In this embodiment, memory 116 also includes an interface identification (I/F ID) look-up table 120. Interface ID look-up table 120 may include, for example, information specifying addresses of interfaces facilitating 20 communication between a core network element 14 and various peripheral network elements 18-24.

The illustrated embodiment of memory 116 also includes a path identification (Path ID) look-up table 122. Path ID look-up table 122 may include, for example, information to 25 facilitate connection oriented routing of, for example, MPLS signals over a specified path to a particular destination network element. In a particular embodiment, path ID look-up table 122 may include information facilitating creation of virtual paths and virtual channels 30 for communication of asynchronous transport mode (ATM), Frame Relay, or packet-over-SONET signals through core cloud 12.

Although look-up tables 118-122 are depicted as separate data structures, one or more of those or other

data structures could be combined without departing from the scope of the invention. In addition, other data structures containing various other information could reside within memory 116.

5 The content and functionality of each core network element 14 may vary depending on that network element's position within core cloud 12 and/or relationship to other network elements. For example, if core network element 14 is coupled to a peripheral network element 18-24, that core network element 14 is likely to include interface look-up tables specifying connection information to various peripheral network elements 18-24.

10
15 In addition, whether peripheral network element 18-24 carries connectionless or connection oriented traffic may effect the content and functionality of the core network element 14 coupled thereto. For example, an intermediate core network element, such as core network element 14n (which couples other core network elements, but not peripheral network elements) typically includes information identifying the topology of core cloud 12, such as Node ID look-up table 118. Where core network element 14 traffics connection oriented signals, memory 116 can also include path identifier look-up table 122. Core network elements 14 may include any combination of these and/or other look-up tables or other data structures, depending on the signals they process and/or the level of processing provided.

20
25
30 Following is an example of various functionality offered by various network elements. Assume a connectionless signal originates at interface equipment 42a and is destined for interface equipment 48a. Interface equipment 42a originates, for example, an Internet Protocol (IP) signal, specifying a network address of destination interface equipment 48a. Interface equipment 42a

communicates the IP signal, including an identification of the network address for interface equipment 48a, across network 34 to router 20. Router 20 communicates the signal across communication link 30d to core network element 14d, 5 the ingress core network element in this case.

Ingress core network element 14d receives the incoming signal from router 20 at peripheral interface 112. Peripheral interface 112 and/or processor 114 examines the incoming signal to determine whether the incoming signal is 10 a connection oriented signal or a connectionless signal, or to determine what type of signaling the signal will experience as it traverses core cloud 12. In this case, the incoming signal is a connectionless IP signal specifying a network address for interface equipment 48a, 15 which is coupled to router 24 through network 38, and further coupled to core network element 14a through communication link 30a. In this example, core network element 14a is an egress core network element for this signal.

Processor 114 appends a transport label to the incoming IP signal, which will direct other core network elements 14 encountering this signal as to how to process this signal. In particular, processor 114 of ingress core network element 14d may create transport label stack 20 comprising a stack of sub-transport labels, each one of the sub-transport labels directing a particular core network element 14 to take a particular action. Each time an intermediate core network element 14 encounters the signal, the core network element 14 examines the sub-transport 25 label on the top of the stack of transport labels and takes an action according to the content of the examined transport label. Once the action associated with the top 30 sub-transport label is taken, that sub-transport label is

popped from the stack, and the next sub-transport label is examined.

The signal is routed from ingress core network element 14d through core cloud 12 to egress core network element 14a in accordance with information specified in the transport label stack appended to the signal. Once the signal reaches egress core network element 14a, processor 114 of egress core network element 14a removes the remaining transport label stack from the signal, and communicates the signal to router 24 through peripheral interface 112. Router 24 may then take appropriate steps to route the signal to its intended destination, in this case, interface equipment 48a.

FIGURE 3A is a block diagram showing exemplary content of a signal 200 for transmission through core cloud 12. In this example, signal 200 includes a signaling data link (SDL) header 210, a transport label 220, a packet area 230, and error detection (ED) information 240.

Signaling data link header information 210 is used to for packet delineation purposes. Signaling data link header 210 may include, for example, an indication of the packet length, an amount of padding added to the packet so that the packet ends on a convenient bit boundary, and Header Error Correction (HEC) information.

Transport label 220 includes an indication of the type of signal being transported, as well as information used in processing the signal during transport. Additional details regarding an exemplary embodiment of transport label 220 will be described below with reference to FIGURE 3B.

Packet area 230 includes protocol information, such as Point-to-Point Protocol (PPP) bits from the original packet header received at ingress core network element 14, and packet bytes copied from the original packet. Packet area 230 may also include padding bytes added to the packet so

that the packet ends on a convenient bit boundary, such as a 64-bit boundary.

5 Error detection information 240 may include various information useful in confirming the integrity of packet payload 230 after its transport through one or more core network elements 14. In a particular embodiment, error detection information 240 may comprise, for example, Cyclic Redundancy Check (CRC)-32 information. Other error correction algorithms could be used consistent with the
10 present invention.

15 FIGURE 3B is a block diagram of exemplary content of transport label 220. In the illustrated embodiment, transport label 220 includes format bits (FMT) 222, a label value 224, quality of service bits (Qos) 226, an end of stack (EOS) indicator 228, and time-to-live (TTL) information 229. In this embodiment, format information 222 includes a series of bits indicating the type of signal being transmitted. In this case, format information 222 comprises an index to a look up table 300 (which will be
20 described with reference to FIGURE 4 below). The information in look up table 300 tells the analyzing core network element 14 what type of signal is being transmitted, and may also include procedures for transmitting that signal.

25 Information within label value field 224 can vary depending on the type of signal being transmitted. Where, for example, signal 200 comprises an IP packet - a connectionless signal - label value field 224 may include a destination network address. Network elements 14 receiving that signal 200 can use the node address contained in label value field 224 to route signal 200 to the next hop. In other cases, where, for example, signal 200 comprises a connection oriented signal such as an MPLS signal with ATM, Frame Relay, or packet-over-SONET

encoding, label value field 224 may include a path identifier. The path identifier could be used to access look up tables to establish, for example, virtual channels and virtual paths for routing connection oriented signal 200. In a particular embodiment, system 10 may receive an MPLS label at the ingress core network element 14, and simply append a format field to the existing MPLS label without further alteration. In that case, the network elements 14 within core cloud 12 process the MPLS label as they would with any other MPLS signal.

Quality of service information 226 may include various information specifying, for example, a priority associated with the particular signal relative to other signals being transmitted by system 10. As a particular example, quality of service information 226 may comprise 3 bits operable to specify one of a variety of transmission priorities.

End of stack indicator 228, in this embodiment, comprises a single bit indicating whether the current transport label is the last transport label associated with this signal 200. In some cases, ingress core network element 14 may generate multiple sub transport labels each associated with the same signal 200. In that case, each sub transport label 220 provides instruction for one or more particular core network elements 14 along the path through core cloud 12 to the destination network element. As core network elements 14 process signal 200 according to each sub transport label, those sub transport labels can be removed from the stack of sub transport labels accompanying that signal 200. The stack of sub transport labels can be designed so that the final sub transport label in the stack directs an egress core network element 14 in its transmission of signal 200 to an associated peripheral network element 18-24. In the illustrated embodiment, end of stack indicator 228 identifies a condition where the

current sub transport label is the last sub transport label associated with signal 200.

Time-to-live field 229 includes information designed to avoid endless looping of information within core cloud 12, and may also be used to limit the scope of a broadcast packet. In a particular embodiment, time-to-live information 229 carries a particular value, which is decremented by one each time a packet passes through one of core network elements 14. If time-to-live value 229 reaches 0 before the packet has reached its destination, the packet can be discarded. Other methods of implementing the concept of a time-to-live variable can be used without departing from the scope of this invention.

FIGURE 4 is a block diagram of an exemplary format look-up table 300. Format look-up table 300 comprises a format index column 310 and an instruction column 312. In the illustrated embodiment, format index column 310 comprises a four bit indicator. Other content and/or format of signal format indicators could be used without departing from the scope of the invention. Core network elements 14 may include a format look-up table such as format look-up table 300 for use in determining what action to take on signals received. In other embodiments, all or a portion of the content of format look-up table 300 could be programmed and/or hard wired into one or more core network elements 14 to facilitate particular processing functions upon recognition of particular information in the format bits 222 of transport labels 220.

In this embodiment, format look-up table 300 includes several "reserved" entries. The functions associated with those entries can be determined, for example, by a network administrator. Other entries 310 are used to specify information associated with a particular type of signal or signal processing. For example, entry 310a (0000)

indicates an MPLS label. Network elements receiving format field 310a will process the information in label value field 224 as an MPLS label. For example, in that case the information in label value field 224 may comprise a path 5 identifier. Core network element 14 analyzing this signal can use the MPLS label in label value field 224 to access path ID look-up table 122 to facilitate generation of a virtual circuit and/or virtual path for this signal.

Entry 310b (0001) indicates that the signal has been 10 designated as a multicast signal. In that case, label value field 224 of that signal will include information facilitating creation of virtual circuits and/or virtual paths to all applicable destination core network elements 14.

Entry 310c (0010) instructs core network element 14 to 15 read the associated label value 224 as an interface identifier specifying a particular interface within an egress core network element 14 to a particular peripheral network element 18-24. The egress core network element 14 receiving this signal can use the value in label-value 20 field 224 to access interface ID look-up table 120 to determine the appropriate interface for that signal. In this example, entry 310c indicates that the associated label value field 224 holds a 32-bit value. Entry 310d 25 (0011) indicates the same information as entry 310c, but in this case associated label value 224 is a 64 bit entry, carrying the interface ID and information regarding the source of the signal. This source information can be useful, for example, for signal control purposes.

Entry 310p instructs core network element 14 to read 30 the associated label value 224 as a node identifier specifying the address of a particular core network element 14. The core network element 14 receiving this signal can use the value in label-value field 224 to access node ID

look-up table 118 to determine the next hop for the signal. This type of label value is typically used for transmitting signals in connectionless mode.

FIGURE 5 is a block diagram showing an exemplary signal format of signal 200 constructed according to the teachings of the present invention. As discussed with respect to FIGURE 3A, signal 200 may include SDL header 210, transport label 220, packet area 230, and packet error detection 240. In this particular example, SDL header 210 comprises 32 bits. The first 13 bits of SDL header 210 comprise an indication of packet length 212, in this case, the number of 32-bit words in the packet. The next 3 bits of SDL header 210 are an indication of a padding length 214. Padding length 214 indicates a number of bits used as padding for the packet data to ensure that the packet area ends on a convenient bit-boundary, for example, a 64-bit boundary. The final 16 bits of SDL header 210 comprise header error correction (HEC) bits 216. The header error correction information comprises Cyclic Redundancy Check (CRC) - 16 information generated over the SDL packet length and padding length fields.

The specific content and format of the information described in SDL header 210 is for exemplary purposes only. Other content and formats could be used without departing from the scope of the invention.

In the illustrated embodiment, transport label 220 of signal 200 comprises a stack of sub-transport labels 220a-220n. As described with respect to FIGURE 3B, each transport label 220a-220n includes format bits 222, label values 224, a quality of service field 226, and an end of stack indicator 228. In the illustrated embodiment, first sub-transport label 220a includes time-to-live field 229. In this embodiment, second sub-transport label 220b includes congestion information 227. Congestion

information 227 may include, for example, a Forward Congestion Notification (FCN), a Backward Congestion Notification (BCN), and a Congestion Experienced indicator (CE), all of which denote that congestion has been 5 experienced by the network. Congestion information 227 may also include an Explicit Congestion Notification Capable Transport indicator (ECT), which denotes that the flow is capable of responding to the congestion experienced indicators. Finally, congestion information 227 may 10 include a Deflectively Routed indicator (DEF), which notes that a packet has been deflected to an alternative path as a protection mechanism. Second sub-transport label 220b may also include a drop precedence 225, which defines the current packet's discard precedence relative to other 15 packets being transmitted.

In the illustrated embodiment, congestion information 227 and drop precedence field 225 reside in the second sub-transport label of the stack of sub-transport labels 220. Each time a sub-transport label 220a-220n is removed 20 from the stack, time-to-live information 229 is copied to the new top sub-transport label 220b-220n, and congestion information 227 and drop precedence information 225 are copied to the current second sub-transport label 220c-220n.

In the illustrated embodiment, each sub-transport 25 label stack 220 will include at least two sub-transport labels 220a-220n. The top sub-transport label 220a will include an indication of the signal type or signal processing type for signal 200. This sub-transport label 220a will instruct core network elements 14 regarding 30 signaling procedures associated with signal 200. As signal 200 traverses core cloud 12 and network elements execute instructions associated with the sub-transport labels 220a-n-1, the sub-transport labels associated with the executed instructions are removed from transport label

stack 220. Core network element 14 then examines the next sub-transport label to receive further instructions on processing signal 200.

In this embodiment, the final sub-transport label 220n in the transport label stack 220 carries an index to interface ID look-up table 120. When the egress core network element 14 receives signal 200 and examines the last sub-transport label 220n, it uses the interface ID in label value field 224n to index its interface ID look-up table 122 and identify an interface between the egress core network element 14 and the destination peripheral network element 18-24. Egress core network element 14 then removes transport label 220n from signal 200 and routes the signal to the destination peripheral network element 18-24.

Packet area 230 contains information as described with respect to FIGURE 3A. In addition, packet area 30 of FIGURE 5 includes ingress node ID information 232 and ingress interface ID information 234, which can be used, for example, in various signal control operations.

In this embodiment, signal 200 also includes error detection information 240, as described in FIGURE 3A.

FIGURES 6A-6E are block diagrams of exemplary signals 200a-200e, and will be used to provide examples of signal processing within core cloud 12. FIGURE 6A-6C assume that an IP signal is received at core cloud 12 from router 24, that the signal is destined for router 20, and that the signal is to be communicated from router 24 to router 20 in a connectionless manner. In these examples, core network element 14a, which is coupled to router 24, serves as the ingress node within core cloud 12, and core network element 14d, which is coupled to router 20 serves as the egress core network element within core cloud 12.

FIGURE 6A is an example of pure connectionless routing of a signal 200a from ingress core network element 14a to

egress core network element 14d. Ingress core network element 14a may format the incoming packet to append transport label 220, or may receive a formatted packet including transport label 200 from another network element. 5 Transport label 220, in this example, comprises a stack of sub-transport labels 220a and 220n. Sub-transport label 220a includes label value 224a comprising a node identification associated with egress core network element 14d. Sub-transport level 220n includes label value 224n comprising an identification of the interface ID for the 10 interface between egress core network element 14d and router 20.

Based on a routing algorithm and the node ID in label value field 224a of sub-transport label 220a, ingress core network element 14a determines the next hop for signal 200a. For example, network element 14a may determine that the most appropriate next hop for signal 200a is core network element 14b. Core network element 14b receives signal 200a and examines the first sub-transport label 220a. 15 In particular, core network element 14b examines format field 222a of sub-transport label 220a to determine that label value field 224a of transport label 220a represents a node ID. Core network element 14b compares its node ID to the node ID specified in label value field 224a. 20 Finding no match, core network element 14b consults its node ID look-up table 118 to determine the appropriate 25 next hop.

In this particular example, network element 14b may determine that the appropriate next hop is to network element 14c. Network element 14c receives signal 200a and, 30 recognizing that the first sub-transport label 220a identifies a node ID, and that network element 14c's node ID does not match the node ID specified in label value field 224a, network element 14c determines the appropriate

next hop for signal 200a. Using the node ID provided in label value 224a, network element 14c consults its node ID look-up table 118 to determine the next hop for signal 200a. In this case, network element 14c communicates 5 signal 200a to egress network element 14d.

Egress core network element 14d receives signal 200a and examines the first sub-transport label 220a. Recognizing, from format field 222a, that label value 224a represents a node ID, egress core network element 14d compares its node ID to the node ID specified in label 10 value 224a. Finding a match, egress core node 14d removes sub-transport label 220a from transport label stack 220 and examines the next sub-transport label 220n. In particular, egress core network element 14d examines format field 222n 15 and determines that label value field 224n comprises an interface ID. Egress core network element 14d uses that interface ID to cross-reference its interface ID look-up table 120 and determine the appropriate interface for transmission of signal 200a to router 20. Egress core 20 network element 14d next removes transport label 220 from signal 200a, and transmits the IP packet to the destination router 20.

FIGURE 6B is a block diagram of a signal 200b 25 formatted for transmission through core cloud 12 using loosely explicit routing. The term "loosely explicit routing" denotes a connectionless form of routing, which specifies one or more intermediate nodes between the ingress node and the egress node of core cloud 12. Following from the previous example, the ingress node in 30 this case is network element 14a, while the egress node in this example is network element 14d. In the example shown in FIGURE 6B, loosely explicit routing is used to specify that signal 200b should traverse through network element 14b on its way to its destination network element 14d. As

a result, transport label stack 220 includes an additional sub-transport label 220b specifying transmission through network element 14b.

In operation, ingress core network element 14a compares the node ID specified in label value 224a, finds no match with its node ID, and consults its node ID look-up table 118 to determine the next hop for signal 200b on its way to network element 14b. In this case, network element 14a communicates signal 200b directly to network element 14b. Network element 14b receives signal 200b, examines format field 222a, and determines that label value field 224a specifies a node ID. Network element 14b compares its node identification with the node identification specified in label value field 224a and determines a match.

Recognizing that the end of stack indicator 228a shows that additional sub-transport labels remain, network element 14b pops sub-transport 220a from transport label stack 220 and continues by examining sub-transport label 220b. Examining format field 222b of sub-transport label 220b, network element 14b recognizes label value field 224b as a node ID. After confirming that the node ID specified in label value field 224b does not match the node ID for network element 14b, network element 14b consults its node ID look-up table 118 to determine the next hop for signal 200b on its way to network element 14d. In this example, network element 14c is identified as the next hop for signal 200b. Network element 14c receives signal 200b and examines the first sub-transport label in transport label stack 220, in this case sub-transport label 220b. Network element 14c ultimately routes signal 200b to egress network element 14d, where signal 200b is routed through the appropriate interface to router 20 in a manner similar to that described with respect to FIGURE 6A.

5 FIGURE 6C provides an example of a method of explicit routing of a signal through core cloud 12. The term "explicit routing" refers to a routing algorithm where each node that will be traversed by the signal through core cloud 12 from the ingress node to the egress node is explicitly specified. To this end, transport label stack 220 comprises a stack of sub-transport labels, each identifying a node ID of a network element 14 in a sequential path between the ingress node and egress node.

10 The final sub-transport label 220n in transport label stack 220 identifies an interface between the egress network element 14 and the destination peripheral network element, in this case, router 20.

15 The example shown in FIGURE 6C assumes that signal 200c is to be transmitted from an ingress network element 14a to an egress network element 14d, by way of intermediate network elements 14n and 14e, sequentially. In operation, ingress network element 14a consults its node ID look-up table 118 to determine an appropriate hop to network element 14n specified in sub-transport label 220a. Network element 14n receives signal 200c and examines format field 222a to determine that label value field 224a is a node ID. Network element 14n compares its node ID to the node ID specified in label value field 224a, and 20 determines a match. Network element 14n then pops the top sub-transport label 224a, and examines the next sub-transport label 224b. Network element 14n determines that label value 224b identifies a node ID, and compares its Node ID to that one. Finding no match, network element 25 14n uses the node ID in label value 224b to index its node ID look-up table 118 and determine a next hop for signal 30 200c, in this case, core network element 14e.

Core network element 14e receives signal 200c, examines format field 222b to determine that label value

field 224b contains a node ID, and compares its node ID to the value in label value field 224b. Finding a match, core network element 14e pops sub-transport label 220b and examines sub transport label 220c. Network element 14e 5 determines that label value 224c identifies a node ID, and compares its Node ID to that one. Finding no match, network element 14e uses the node ID in label value 224c to index its node ID look-up table 118 and determine a next hop for signal 200c, in this case, core network element 10 14d.

When signal 200c arrives at egress core network element 14d, the process continues similarly to that described with respect to FIGURE 6A.

FIGURE 6D shows a block diagram of a signal 200d, 15 which contains information associated with a connection oriented signal. This example assumes that interface equipment 40a desires to communicate an MPLS signal with ATM encoding to interface equipment 44a. In this example, core network element 14b coupled to label switching router 18 is an ingress node, and core network element 14e coupled 20 to label switching router 22 is an egress node on core cloud 12. The example assumes that signal 200d is to be routed through core cloud 12 using MPLS routing.

Ingress core network element 14b creates, or receives 25 signal 200d with transport label stack 220. In this example, transport label stack 220 includes sub-transport label 220a, which specifies in label value 224a a path identifier associated with signal 200d. Transport label stack 220 also includes sub-transport label 224n, which 30 specifies an interface address between egress core network element 14e and label switching router 22. Ingress core network element examines format field 222a of signal 200d, and determines that label value 224a specifies a path ID. Using the MPLS label in label value field 224a to index its

path ID look-up table 122, ingress core network element 14b establishes virtual circuits and/or virtual paths for communicating signal 200d to egress core network element 14e.

5 Egress core network element 14e receives signal 200d, pops the top sub-transport label 224a, and examines sub-transport label 224n to identify the interface linking that core network element 14e and label switching router 22. Core network element 14e restores signal 200d to its
10 original format, and communicates signal 200d to label switching router 22 for further processing.

FIGURE 6E is a block diagram of a signal 200e, which will be routed through core cloud 12 using both connectionless and connection oriented signaling. This example assumes a signal is received from label switching router 18, and is to be communicated through core cloud 12 to label switching router 22. This example also assumes that the signal is intended to be transmitted through core cloud 12 in a connectionless manner until it reaches core network element 14d, and then transmitted to egress core network element 14e using a connection oriented process. This type of signal routing could be desirable, for example, where particular regions of core cloud 12 require traffic engineering capabilities due to their relative
20 congestion compared to other areas of core cloud 12.
25

To accomplish this hybrid routing process, system 10 generates transport label 220, which includes sub-transport label 220a – specifying a node ID for connectionless signaling to core network element 14d; and sub-transport label 220b – specifying a path ID from core network element 14d to egress core network element 14e to facilitate connection oriented signaling during that part of the signal path.

In this example, ingress core network element 14b generates or receives transport label stack 220 associated with signal 200e. Network element 14b examines sub-transport label 220a and determines that transport value 224a comprises a node ID that does not match the node ID for ingress core network element 14b. Using transport value 224a to index its node ID look-up table 118, core network element 14b identifies the next hop for signal 200d.

This connectionless signaling method continues as described above with respect to FIGURES 6A-6C, until signal 200d reaches core network element 14d. At that point core network element finds a match between its node ID and the node ID specified in label value field 224a, removes sub-transport label 220a, and examines the next sub-transport label 220b. Finding that label value 224b includes a path ID, core network element 14d uses the label value 224b to index its path ID look-up table, and facilitates connection oriented signaling of signal 200d from core network element 14d to core network element 14e. At core network element 14e, sub-transport label 220b is popped, and sub-transport label 220n is used to communicate the packet within signal 200d to an interface linking egress core network element 14e and label switching router 22.

FIGURE 7 is a flow chart showing an exemplary method 300 of communicating connectionless and connection oriented signals over common network elements. The method 300 begins at step 310 where core cloud 12 receives a signal from a peripheral network element 18 - 24, and formats the signal for transport through core cloud 12. Method 310 may include receiving a connectionless signal, a connection oriented signal, or a signal requiring both connectionless and connection oriented processing at step 312. The

ingress core network element 14 receiving the signal from peripheral network elements 18 - 24 determines at step 314 the signal type or the type of routing associated with that signal. Ingress core network element 14 appends transport 5 label 220 to the signal at step 316. Ingress core network element 14 may generate transport label 220, or may receive the transport label from another network element. In a particular embodiment, transport label 220 comprises a stack of sub-transport labels, each sub-transport label 10 specifying a routing instruction for all or a portion of the signal's path through a core cloud 12.

The method continues at step 320 where system 10 transports formatted signal 200 through core cloud 12 to an egress core network element 14. Details regarding an exemplary method of transporting signal 200 through core 15 cloud 12 will be described below. Egress core network element 14 receives formatted signal 200 and returns formatted signal 200 to its original format at step 360. In a particular embodiment, egress core network element 14 20 returns signal 200 to its original format by removing transport label 220 from signal 200. Egress core network element 14 communicates the original signal to the destination peripheral network element 18-24 for at step 370.

25 The following description provides additional details on an exemplary method 320 of transporting signal 200 through core cloud 212. The method 320 begins at step 322 where core network element 14 receiving the formatted signal examines format field 222 of transport label 220. 30 If core network element 14 determines at step 324 that label value 224 comprises a node ID, core network element 14 determines at step 326 whether its node ID matches the node ID specified in label value field 224. If the node IDs match, core network element 14 pops the top

sub-transport label 220a from transport label stack 220 at step 328, and examines the next sub-transport label 220b at step 322.

5 If format field 222 identifies label value 224 as a node ID at step 324, and core network element 14 determines at step 326 that its node ID does not match the node ID specified in label value field 224, core network element 14 determines at step 330 the next hop for signal 200.
10 Determining the next hop for signal 200 may involve, for example, using the node ID specified in label value field 224 as an index to node ID look-up table 118, and applying a routing algorithm accordingly. Once core network element 14 determines the next hop for signal 200, it communicates
15 signal 200 to the next core network element 14 at step 332.

After examining format field 222 of transport label 220 at step 322, core network element 14 may determine at step 334 that label value 224 is an MPLS label containing, for example, a path identifier. In that case, core network element 14 uses label value 224 to access, for example, path ID look-up table 122 to facilitate generation of virtual paths and/or virtual circuits for connection-oriented communication of signal 200. In some cases, label value 224 may comprise a multi-cast label. In those situations, core network element 14 may generate virtual circuits and/or virtual paths for each destination core network element at step 336. Destination core network element 14 receives signal 200 at step 338, and may pop the top sub-transport label 220a-220n-1 at step 340. The destination core network element 14 then examines format field 222 of the next sub-transport label at step 332.

Various signals may use some or all of steps 322-340 described above. Some signals may use purely connectionless signaling, some may use only connection

oriented signaling, and others may use a combination of connectionless and connection oriented signaling.

When core network element 14 examines format field 222 of transport label 220 at step 322 and determines at step 5 342 that label value 224 of the current sub-transport label comprises an interface ID, core network element 14 determines the correct interface between it and the destination peripheral network element 18-24. For example, egress network element 14 may use the interface ID as an 10 index to interface ID look-up table 120 to identify the correct interface between that core network element 14 and the destination peripheral network element 18-24. Egress core network element 14 returns signal 200 to its original format at step 360 and communicates the original signal to 15 the destination peripheral network element 18-24 at step 370.

At any time during the signaling process, core network element 14 may examine format field 222 of transport label 220 and determine that label value 224 does not specify a 20 node ID, a path ID, or an interface ID. Rather, in those cases, label value may specify a particular processing function that may be defined, for example, by a network administrator. In those cases, core network element 14 performs the specified signal processing functions at step 25 344.

Although the present invention has been described in several embodiments, a myriad of changes, variations, alterations, transformations, and modifications may be suggested to one skilled in the art, and it is intended 30 that the present invention encompass such changes, variations, alterations, transformations, and modifications as fall within the spirit and scope of the appended claims.